END
FILMED
DTIC

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A145 711

(12)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS
BEFORE COMPLETING FORM

1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
Technical Report #3		

4. TITLE (and Subtitle)	5. TYPE OF REPORT & PERIOD COVERED
Toward a Computational Neuropsychology of High-Level Vision	Interim Report
	6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s)	8. CONTRACT OR GRANT NUMBER(s)
Stephen M. Kosslyn	N00014-82-C-0166 N00014-83-K-0095

9. PERFORMING ORGANIZATION NAME AND ADDRESS	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
Department of Psychology & Social Relations Harvard University Cambridge, MA 02138	NR 150-480

11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE
Personnel & Training Research Programs Office for Naval Research Arlington, VA 22217	August 20, 1984
	13. NUMBER OF PAGES

14. MONITORING AGENCY NAME & ADDRESS(If different from Controlling Office)	15. SECURITY CLASS. (of this report)
	UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release: distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

DTIC
SELECTE
SEP 11 1984
E

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Cognitive science   visual processing
neuropsychology   artificial intelligence
mental imagery
computer models

DTIC FILE COPY

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Visual information processing in humans has recently been studied from three distinct perspectives, with only minimal cross-fertilization among them. This paper develops a way of melding the approaches of Artificial Intelligence, Cognitive Psychology, and Neuropsychology, and explores the advantages of such a hybrid approach. Each of the individual approaches has its strengths and weaknesses, but these are different for the

DD FORM 1473   EDITION OF 1 NOV 65 IS OBSOLETE
1 JAN 73
S/N 0102-LF-014-6601

84   09   06   047

different approaches, it is argued in this paper that by combining the three, we are in a position to take advantage of each one's strengths and may be able to circumvent each one's weaknesses. Additional keyword topics include mental imagery and computer models.

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

TOWARD A COMPUTATIONAL NEUROPSYCHOLOGY OF HIGH-LEVEL VISION

Stephen Michael Kosslyn

Visual processes in humans have recently been studied from three distinct perspectives, with only the barest amount of cross-fertilization among them. In this chapter we consider a way of melding the approaches of Artificial Intelligence (AI), Cognitive Psychology and Neuropsychology, and explore the advantages of such a hybrid approach. Each of the individual approaches has its strengths and weaknesses, but these are different for the different approaches; by combining the three, we are in a position to take advantage of each one's strengths and may be able to circumvent each one's weaknesses. Although I believe that most of the observations I will make in this chapter generalize to the study of all cognitive abilities, I will restrict the examples to vision. Vision has been the subject of intense study in the three disciplines, and the evidence seems clear at least in this case that there is much to be gained by combining the approaches.

The focus in this chapter is on just those events that take place near the end of the visual processing sequence that originates at the eyes. These events can be considered "mental" because they can be affected by one's knowledge and beliefs (whereas processes carried out by low-level systems, such as those localized at the retina, presumably are not affected by one's knowledge and beliefs). The study of high-level, "mental" events presents problems that are not as severe when one studies "low-level" processing, which is closely tied to properties of the stimuli. In low-level vision, an analysis of the geometry of surfaces and the optics of light place strong constraints on how information must be processed, as we shall see below. By the time we get

to high-level processing, however, these properties of the stimuli have been transformed numerous times in numerous ways. How can we best go about trying to understand the last phases of the sequence of transformations? This task is a little like constructing a ship at sea, with each piece floating freely. Once we have nailed down a few of the pieces, the job will become easier; but how do we identify those initial pieces? Let us briefly review the key features and limitations of the approaches currently taken in AI, Cognitive Psychology and Neuropsychology.

1. The Computational Approach

One way of trying to understand the nature of vision is to consider what would be necessary to program a computer to see. In so doing, one is first led to ask about the purposes of vision, and then is led to consider what problems must be solved in order for it to serve these ends. At the most general level of analysis, vision serves three functions: First, it allows one to identify objects and events in the environment. Central to this capacity is the ability to compare representations of input to stored representations of previously-seen objects. Second, it allows one to navigate around in the environment (without bumping into objects), and conversely, to avoid or intersect other objects that are moving. Central to this capacity is the ability to represent metric spatial relations and to update them efficiently as the organism or part of the environment moves. Third, it allows one to reason about objects and events in their absence (e.g., to consider whether one's hand could fit into a certain hole one remembers being of a specific size and shape). Central to this capacity is the ability to "re-present" objects and events to oneself in their immediate absence and to operate on these

representations in a way that will allow one to anticipate what would happen

should the analogous actual operations be performed in the real world.

In trying to understand even one of these capacities, researchers very

quickly discovered the usefulness of positing a modular design, with separate

mechanisms being used to carry out distinct aspects of performance. Thus,

researchers in AI developed theories of the processing modules used in

vision. A processing module is a "black box" that carries out specific

computation or computations. By "computation" I mean, roughly, "a meaningful

(i.e., informationally interpretable) transformation of an input." The

theorist specifies the nature of the computations performed by various modules.

A theory of a computation specifies three things: the information

available to be used in performing a computation, the purpose of the

computation, and a description of what is being computed (see Marr, 1982). For

example, consider a theory of a computation used in low-level vision to detect

edges of objects. The information available is an intensity array, with

intensity values specified for each point on the image. The purpose of the

computation is to discover places where the intensity changes rapidly, which

are assumed to correspond to edges of represented surfaces. What this

computation does can be described as finding the zero-crossings in the second

derivative of the function relating intensity and position. (The actual theory

is more complicated, involving a convolution of the image with a function

representing the output from the very early processors; however, this brief

presentation is sufficient for present purposes. See chapter 3, Marr, 1982).

Researchers in AI do not stop with theories of processing modules and

their constituent computations. Rather, in order actually to build a working

program one must also formulate a theory of how a computation is actually accomplished on-line. Each "black box" can be opened up, so to speak, and its internal workings described. Indeed, a theory of processing modules (and their associated computations) is a way of organizing sets of representations and processing operations into coherent units. That is, a processing module is presumed to correspond to a mechanism that accomplishes the computations that constitute the module. The on-line operation of this mechanism can be described by a theory of the algorithm for a given task. The algorithm specifies step by step how a computation is carried out.

To get a feel for the distinction between a computation and the algorithm that carries it out, think of the number of different ways one could perform a computation like multiplication; one could add one of the numbers to itself over and over, convert the numbers to logs and add the exponents, etc. The actual procedure follows an algorithm, and numerous different algorithms can be used to carry out the same computation.

In the course of developing theories of the algorithms used, a theory of the functional architecture is developed. (Newell & Simon, 1972, are primarily responsible for introducing the idea of a functional architecture to psychology.) A theory of the functional architecture specifies the kinds of representations (e.g., Roman numerals, numbers in log base 10, etc), buffers (places where representations can be stored), and processing operations (such as addition, matching, and substitution) that can be used in the algorithms that actually carry out the computations (see Kosslyn, 1984, for a more detailed discussion of the concept of a functional architecture). A given component of the functional architecture (e.g., a buffer) in principle could be

used by different algorithms that carry out different computations (e.g., the same buffer can be used to store two numbers being added or multiplied), or it could be used only by one algorithm, which carries out only a single computation.

A "computational theory," then, is a theory that 1) specifies the processing modules (and the constituent computations) used in performing a set of tasks; 2) specifies the representations, buffers, and processing operations used in carrying out the computations; and, 3) specifies the precise sequence of steps used to perform a set of tasks. Incomplete computational theories are today the rule rather than the exception, but all computational theories are directed at eventually specifying these three aspects of information processing.

Limitations of the approach

On Marr's view, the core of a theory of how information is processed is the theory of the computation. The notion of a theory of the computation is relatively novel for cognitive psychology, and it is worth exploring the force of Marr's views. Marr (1982) argues that the information available and the purpose of a computation often virtually dictate what the computation must be. This sort of theory can sometimes be almost like a solution to a mathematics problem, arising through logical analysis of the nature of the problem to be solved and the input available to solve it. That is, if the task is very well defined, and the input is highly restricted, a specific computation may almost be logically necessary. Further, Marr claims that once a computation is defined the task of characterizing the representations and processes used in carrying out an algorithm is now highly constrained: the representation of the

input and the output must make explicit the information necessary for the computation to serve its purpose (e.g., picking out likely locations of edges), and the representations must be sensitive to the necessary distinctions, be stable over irrelevant distinctions, and have a number of other properties (see Marr, 1982, chapter 5).

To return to the example of the computation for detecting edges that was discussed above, note that once we have described the purpose and the input, we have almost defined what has to be computed. In addition, once the theory of the zero-crossings computation was formulated, the theory of the representation of the ouput of the computation was highly constrained: it needed to have primitives that were likely to correspond to physically-meaningful properties of the geometry of surfaces, and had to make explicit places where zero-crossings exist. Marr's "primal sketch" uses short line segments, bars, blobs and the like to connect contiguous zero-crossings, producing a representation with properties that are desirable as input to later computations that derive characteristics of surfaces and shape.

Marr's strong claims about the priority of the theory of the computation do seem appropriate for some of the problems of low-level vision, but only because there are such severe constraints on the input (posed by the nature of the world and the geometry of surfaces) and because the purpose of a computation is so well-defined (e.g., to detect places where intensity changes rapidly, to derive depth from disparities in the images striking each eye, to recover structure from information about changes on a surfaces as an object moves). In cognition, the situation is somewhat different: First, the basic abilities in need of explanation--analogous to our ability to see edges or to

see depth in vision--must be discovered. For example, with the advent of new
methodologies, our picture of what can be accomplished in mental imagery has
changed drastically (e.g., see Shepard & Cooper, 1982). Second, the input to a
"mental" computation often is not obvious, not necessarily being constrained by
some easily-observed property of the stimulus. One must have a theory of what
is represented before one can even begin to specify the input to the
computations. Third, the optimal computation will depend in part on the kinds
of processing operations that are available; presumably, over the course of
evolution new computations developed in part by taking advantage of the
available processing resources. Thus, developing a theory of the functional
architecture--which specifies the types of representations and processing
operations available--would seem to go hand in hand with developing a theory of
a cognitive computation.

This conclusion is illustrated by problems with some of Marr's own work on
"higher level" vision. Marr posits that shapes must be stored using
"object-centered" descriptions, as opposed to "viewer-centered" descriptions.
In an object-centered description an object is described relative to itself,
not from a particular point of view. Thus terms such as "dorsal" and "ventral"
would be used in an object-centered description, as opposed to terms such as
"top" and "bottom" which would be used in a viewer-centered description. Marr
argues that because objects are seen from so many different points of view, it
would be difficult to recognize an object by matching viewer-centered
descriptions of input to stored representations. However, this argument, based
on a theory of the purpose of the computation, rests on implicit assumptions
about the kinds of representations and processes available in the functional

architecture. If there is an "orientation normalization" pre-processor, the argument is obviated: in this case, a viewer-centered description could be normalized (e.g., so the longest axis is always vertical) before matching to stored representations. And in fact, we do "mentally rotate" objects to a standard orientation when subtle judgments must be made (see Shepard & Cooper, 1982). Further, the mere fact that we do seem to normalize the represented orientation, at least in some cases, casts doubt on the power or generality of object-centered representations (if object-centered descriptions are made, it simply is not clear why orientation normalization would be necessary). In fact, when the matter was put to empirical test, Jolicoeur & Kosslyn (1983) found that people can use both viewer-centered and object-centered coordinate systems in storing information, and seem to encode a viewer-centered one even when they also encode an object-centered one, but not vice versa.

Similarly, arguments can be levied against Marr's assumption that the representations are genuine 3 dimensional representations, as opposed to "2 1/2-D" representations, where one only stores the visible depth information (and not the occluded parts, as opposed to an actual 3-D representation, which stores all parts--as would occur in a stick figure or pattern of points in a 3-D array). Further, one can even question whether shape representations used in recognition are distinct from those used in navigation and visual reasoning .as is involved in deciding whether a jar can fit on a particular space in the refridgerator). If not, then the input to the recognition computation is apt to be quite different from what was assumed by Marr.

The point is that a logical analysis of requirements on the computation is not enough: at least for high-level abilities, the specifics of a computation

will depend to some extent on what types of representations and processing
operations are available in the functional architecture. One can only discover
the actual state of affairs empirically, by actually studying the way the brain
works.

Although the computational approach is not sufficient in and of itself to
lead one to formulate a correct theory of information processing, it does have
a lot to contribute to the enterprise: Thinking about how one could build a
computer program to emulate a human ability is a very useful way of enumerating
alternative processing modules, functional architectures, and algorithms. Not
only does this approach raise alternatives that one many not have otherwise
considered, but it eliminates others by forcing one to work them out concretely
enough to reveal their flaws (the Guzman approach to vision is a good example;
see Winston, 1975).

II. The Cognitive Psychology Approach

The approach in cognitive psychology has been solidly empirical.
Researchers have developed methodologies that make use of response times, error
rates and various judgments, and have developed ways of using these
methodologies to draw inferences about underlying mechanisms. The
methodologies used have become very sophisticated and powerful, allowing
researchers to observe quite subtle regularities in processing. As we saw in
the previous section, such data place strong constraints on theories of
processing: since processing takes place in real time, there will always be
measurable consequences of any given sequence of activity--and if the wrong
pattern of responses occurs, a theory can be ruled out.

Although the psychologists occasionally focus on the nature of an

algorithm a subject is using (particularly if the subject is an expert at the

activity, e.g. see Simon & Simon, 1978), they usually have been interested in

studying specific components of the functional architecture (e.g., a short-term

memory buffer; organization of a long-term memory network; types of production

rules). Properties of components of the functional architecture are revealed

when a person is engaged in a specific kind of information processing that

presumably requires use of those components. However, it has proven difficult

to draw firm conclusions about the underlying architecture or algorithms

because of two general problems: structure/process tradeoffs and task demand

artifacts.

Structure/process tradeoffs

Anderson (1978) demonstrated that given any set of data, more than one

theory can always be formulated to account for the data. His proof rests on

the pervasive possibility of "structure/process tradeoffs." That is, what in

one theory are properties of a given representation operated on by a specific

process are in another theory properties of a different representation operated

on by a different process (and this process compensates for the difference in

representations, producing the same input/output characteristics when the

representation is operated upon). The "analogue/propositional" imagery debate

provides a good illustration of this point. For example, consider the results

of experiments on "mental rotation" (see Shepard & Cooper, 1983, for a review),

in which subjects require increasingly more time to compare two similar figures

that are presented at increasingly disparate orientations. The "analogue

theories" posit a representation that depicts the objects. That is, 1)

each part of the representation corresponds to part of a stimulus such that, 2)

the distances among parts in the representation (where "distance" is defined functionally--as are distances among cells in an array in a computer) preserve the actual distances among the corresponding parts. These representations are like patterns of points in an array in a computer, and rotation is accomplished by shifting the points incrementally--with more shifts being required to effect a greater change in the represented orientation (see Kosslyn, 1980; 1981).

In contrast, "propositional theories" posit that objects are always represented in terms of descriptions. In this case, each part is described as being in a certain position relative to another part (e.g., attached to the left and oriented 45 degees up), and "rotation" consists of altering the relations incrementally (e.g., changing the number representing the angle from 45 to 90 degrees in 15 degree steps). Thus, greater "rotations" require more time.

The two types of theories mimic each other, but in a rather uninteresting way: they are created ad hoc simply to account for the data. What is required are constraints on the theories, a source of motivation for selection of the specific representations and processes. Why should information be represented depictively or propositionally? Why is the transformation apparently done incrementally? Computational considerations are one possible source of constraint (e.g., a depictive representation makes explicit all metric spatial relationships among an object's parts, which is very useful for performing certain kinds of computations). However, we saw above that computational constraints in and of themselves are not sufficient--and in fact the observation of how the system functions (i.e., the dependency of response time on angle) put constraints on computational theories themselves.

Anderson (1978) drew some very pessimistic conclusions from the possiblity of speed/accuracy tradeoffs, but others such as Hayes-Roth (1979) and Pylyshyn (1979) were less gloomy. The upshot of the debate seems to be that it is possible to derive firm inferences about processing mechanisms from behavioral data, but it is very difficult to do so. One argument to be developed in this chapter is that neuropsychological data are powerful supplements to the usual behavioral data, and greatly diminish the ease of using structure/process tradeoffs to concoct alternative theories.

Task demands

Another problem in interpreting behavioral data is the possibility of task demands, which is especially severe in studies of visual thinking. That is, subjects may respond (e.g., by taking longer to rotate an image of an object oriented at a greater angle) because they believe--perhaps unconsciously--that this is what the task requires them to do. Part and parcel of understanding the task may be to mimic the analogous real world event (cf. Pylyshyn, 1981). If so, then data from many studies of mental imagery may say nothing about the nature of the underlying mechanisms, but only reflect the subjects' understanding of tasks, knowledge of physics and perception, and ability to regulate their response times.

Although the problem of task demands has been brought to our attention primarily in the imagery literature (see Kosslyn, Pinker, Smith & Shwartz, 1979, and commentators on that paper), it is applicable to many domains in cognitive psychology. There is no way to ensure that subjects are not unconsciously producing data in accordance with their "tacit knowledge" about perception (and cognition) and their understanding of what the task requires

them to do. In contrast, neurological maladies not only produce behavioral

deficits of various types, but often the patients are not aware of the nature

of these deficits (as will be discussed below for "unilateral visual neglect").

Thus, these types of data might profitably supplement the usual cognitive data

if for no reason other than to rule out task demand accounts of data (to the

extent that patients cannot be responding to task demands because they are

unaware of what they cannot do). And such data are useful for other purposes,

as discussed in the following section.

In short, the strong suit of the cognitive psychologists is their

sophisticated methodologies and the well-described phenomena discovered in the

laboratory. However, although these phenomena can be used to rule out theories

that posit specific structures operated on by specific processes, they are

difficult to use to pin down the properties of specific aspects of the

functional architecture; a theory must explain the data, and although many

cannot, there remain many that can. As will be discussed shortly,

neuropsychological data help to put two important kinds of constraints on

theories of how information is processed: constraints on the nature of the

processing modules, and constraints on the representations and p   essing

operations used in the modules. However, these data are useful only if

construed within a theoretical framework--which can be provided using a

computational approach--and if approached with sensitive methodologies--which

have been developed in cognitive psychology.

III. Neuropsychological approaches

It is important to begin by distinguishing between two related, but

distinct, neuropsychological projects: On the one hand, one can focus on the

theory of functioning. That is, one could use neuropsychological data (e.g., behavioral dysfunction following brain damage) to help formulate and evaluate the computational theory. On the other hand, one can focus on the brain per se. In this case, one would try to characterize different brain loci (e.g., cerebral hemispheres) or patterns of activation in terms of the computations they support. In this chapter the focus is on the theory of functioning, although in developing such a theory we may discover interesting facts about the role of specific brain structures. It is my belief that a good computational theory will provide a good "road map" to guide investigations of the operation of the organ itself, and may even be a necessary prerequisite to understanding how the brain itself works.

The fact that cognition is something the brain does is so obvious it seems barely worth stating. But because of this fact, if a theory of cognitive processing is correct, then the various distinctions made in the theory must be respected by the brain. For example, if a theory claims that shape and color are extracted by separate mechanisms, then separate mechanisms must exist in the brain (which need not be localized in distinct regions, however). The nature of the brain introduces a number of constraints for theories of cognition: The theory should be able to explain why certain abilities are lost together whereas others can be lost separately. It should also be able to explain why patterns of brain activity are more or less similar for different sorts of tasks. Furthermore, theories must obey the broad constraints imposed by the nature of the mechanism itself; for example, if a theory posits that items are searched at a rate exceeding the firing time of neurons, the theory must be incorrect. Thus, it makes sense to look at data that bear on the

functioning of brain mechanisms when formulating and testing theories of
cognition.

Neuropsychological data are of two broad classes: First, and by far the
most predominant, are data on behavioral dysfunction following brain damage.
The damage can be endogenous (e.g., following a stroke or development of a
tumor) or exogenous (e.g., following head injury or surgery, as in split-brain
patients). Second, and of more recent vintage, are data on dynamic changes
within an intact brain performing specific cognitive tasks. These data are
obtained primarily by using EEG (electroencephalographs), ERP (event-related
potentials), PET (positron emission tomography), Xenon-133 regional cerebral
bloodflow, and NMR (nuclear magnetic resonance) techniques. Each technique has
different advantages and drawbacks, and to a large extent they complement one
another.

John Hughlings Jackson is usually credited with making the first
substantive observations on visual deficits following brain damage. In 1874 he
described a way in which the cerebral hemispheres might be specialized,
proposing that the posterior part of the right hemisphere is the "chief seat of
the revival of images in the recognition of object, places, etc. (pg 101)".
This inference was based on the problems a patient with a lesion in this area
had in knowing where she was. In 1876 Jackson described what is now known as
"visual agnosia" (also called "mindblindness"); this patient failed to
recognize her nurses, got lost frequently when travelling familiar routes, and
often did not know objects, persons or places. This malady resulted from a
lesion in the posterior right hemisphere. Patients suffering from visual
agnosia are not blind: these patients can compare two shapes reliably when

both are visible, but they cannot visually recognize what an object is
(although many can recognize objects by touch). This sort of agnosia has been
well-documented in the literature (see Benton, 1982). By 1910 a number of
visual/spatial deficits following brain damage had been identified, including
difficulties in reading, locating objects in space, and "neglect" (ignoring) of
objects that lie off to one side of the viewer. In addition, various theorists
(e.g., Rieger, 1909; Reichardt, 1918--discussed in Benton, 1982) hypothesized
that spatial/practical functions and verbal/conceptual functions are carried
out by distinct mechanisms, which might be localized to the cerebral
hemispheres (with verbal/conceptual on the left, spatial/practical on the
right).

Recent reviews of the literature on visual deficits following brain damage
(e.g., Benton, 1982; Ratcliff, 1982) reveal that we now know that various
clinical signs are fairly common following damage in particular regions of the
brain (e.g., neglect of the left side often follows damage to the right
parietal lobe), and we know that various deficits can be dissociated. For
example, patients can have difficulty in copying objects (by drawing or
constructing a model) but have no visual discrimination problems, or vice versa
(Costa & Vaughan, 1962). In addition, considerable effort has been made in
trying to identify various functions with one hemisphere or the other (see
Springer & Deutsch, 1981).

Perhaps the most important conceptual development in the brain damage
literature is the formulation of the logic of the "double dissociation." If
some behavioral deficit reflects damage to a specific processing mechanism
(e.g., for performing some sort of shape discrimination), and at least part of

this mechanism is distinct from other processing mechanisms, then one should find cases where the ability is spared while other abilities are disrupted (e.g., perhaps discriminating orientation) and vice versa. (It is the "vice versa" that produces the "double" dissociation.) This sort of data provides very strong evidence for a particular configuration of processors in the system.

In addition to dissociations, one also finds associations. If a patient cannot perform task X, in many cases he or she also will be unable to perform task Y. This sort of result could indicate that the same aspect of the functional architecture is recruited in performing both tasks, and that component no longer functions effectively. However, one must be careful here: it could be that different functions happen to be carried out by the same (or nearby) cortical tissue, and hence the association of deficits following brain damage in a given region says nothing about shared functions in different tasks. Thus, careful tests must be devised to ensure that processing is disrupted in the same way in different tasks in order to provide evidence that the tasks share a processing component (I will provide an example of this in the next section).

Limitations of the approach

There are two limitations evident in the neuropsychological literature that are of particular interest here: First, the theories have not been very sophisticated. For example, "localizing oneself in space" is usually considered a single ability in the neuropsychological literature, whereas a computationally-oriented theorist would be inclined to decompose this ability into various encoding, representation and retrieval operations. Similarly,

visual agnosia is described ("mindblindness"), but the underlying causes of the deficit have not be explained; a computational approach would lead one to attempt to characterize the nature of the representations (or properties thereof) that may be lost or to characterize the nature of the failure of processes that encode perceptual information, match it to stored input, and make use of the stored information.

The computational approach has recently had an influence in neuropsychology, and appears to be a promising avenue for future work. For example, Moscovitch (1979) distinguishes between low-level "stimulus features" (presumably processed by both hemispheres) and higher order processes (which may be localized in one cerebral hemisphere). This distinction helps to explain why hemispheric specialization only appears for some phenomena. A more detailed computational analysis might reveal that a given type of stimulus feature (e.g., places where intensity changes most rapidly) might rely on a computation that is localized in a given region, whereas others might rely on computations localized in other regions. Thus, guided by such notions a closer look might reveal subtities that are not evident in the available data. An analogous case is our study of image generation, discussed in the following section, which illustrates how a computational analysis can illuminate neuropsychological phenomena.

The second limitation evident in many neuropsychological studies with humans (but not usually with animals; e.g., see Ungerleider & Mishkin, 1982) is a lack of sophisticated methodologies. Much neuropsychological work centers on administering standardized tests to various patient populations and looking for differences in performance. These tests, however, do not necessarily tap

distinct underlying processing mechanisms, and performance on them may be related in a complicated way to underlying deficits.

In short, neuropsychological data provide another source of constraint on theories of high-level visual processing. They have the potential of being especially useful in identifying processing modules, given the logic of "double dissociation". Let us now consider in more detail some of the potential benefits of combining the three approaches.

IV. Combining the Approaches

The logic of dissociations and associations in deficits is a very powerful way of developing and testing computational theories if it is yoked with the methodologies and analytic techniques developed by the cognitive psychologists. The methodologies developed by the cognitive psychologists for the most part can be adapted for use in neuropsychological studies (much as many of them have been adapted to study cognitive processes in children; e.g., see Siegler, 1978). However, this has not been done by the few researchers who have used neuropsychological data to place constraints on explicit computational theories of high-level vision. For example, Marr (who was perhaps the best computational theorist, and thus worthy of such close examination) was very impressed by the Warrington & Taylor (1973) findings on the failure of patients with parietal-lobe damage to recognize mis-oriented objects (e.g., buckets viewed from the top). Marr concluded that this failure demonstrated that objects were stored as descriptions, and that descriptions were structured around assigning a major axis to an object and then minor axes (of attached parts) off of it; when buckets were seen top down, one presumably had difficulty locating the major axis. Unfortunately, the patient's problems may

have had nothing to do with axis assignment: perhaps they were unable to "mentally rotate" the buckets into a connonical orientation during the recognition process. This possibility is, of course, directly testable by applying the methodologies of contemporary cognitive psychology to brain damaged populations.

To summarize, each of the three approaches discussed above has something to offer, and each is complemented by the other two:

The computational approach is especially useful for generating hypotheses about processing mechanisms: Thinking about the requirements of the task at hand and how one would need to program a computer to perform it is a good way of generating alternative possiblilities. In addition, this approach provides a way of testing complex theories, by actually building a computer program that emulates cognitive processing (see Newell & Simon, 1972). Precise theories of on-line brain functioning may well be so complex that many of a theory's implications will be derived only by using simulation models.

Neuropsychological data offer constraints both on theories of processing modules and theories of the functional architecture. The finding of double dissociations allows one to argue that abilities involve at least some specialized processing modules. In addition, as will be illustrated shortly, the finding of specific deficits that generalize across tasks of a given type can be used to implicate specific representations and buffers. However, neuropsychological data are open to multiple interpretations (just as are any other data), and must be approached analytically.

The methods of cognitive psychology can be profitably used analytically to investigate computational hypotheses about neuropsychological phenomena. These

methods allow one to isolate the variables responsible for an effect, and often specific variables can be identified as reflecting the operation of distinct computations (e.g., using Sternberg's, 1966, additive factors methodology). In addition, once there are prior reasons for positing a specific modular composition of the system, the standard techniques of cognitive psychology become more powerful: Once a module is defined, the number of "degrees of freedom" is reduced for possible structure/process tradeoffs. That is, without modularity constraints, any part of the system can be invoked in combination with any other part to explain a specific result; but if a result can be shown to rest on the operation of a specific module--which is distinct from other modules--the explanation of the result becomes more constrained. Once well-specified classes of alternative theories are defined, cognitive psychologists are better able to specify which phenomena will distinguish among competing accounts (e.g., see the mental rotation case discussed above as treated in chapter 8 of Kosslyn, 1980).

Thus, the three approaches complement each other. The very rich neuropsychological phenomena place strong constraints on computational theories, especially when the tools of cognitive psychology are used to precisely characterize the phenomena. In addition, the computational approach provides useful guidelines about which phenomena are worth detailed scrutiny (as illustrated above in the discussion of Marr's use of Warrington & Taylor's findings). Furthermore, theory development will become much more challenging--and potentially rewarding--if we combine the requirements from all three disciplines: The theory must not only explain the neuropsychological phenomena and the data from normal subjects, but ultimately must be capable of

guiding one to build a computer model that actually emulates the behavior of
normal and brain-damaged subjects. Unlike the case in cognitive psychology,
where it is easy to construct numerous alternative theories, we will be lucky
to formulate even a single theory that meets these criteria.

V. Some Examples of A Computational Neuropsychology
of High-Level Vision

It is probably most useful to provide some concrete examples of how this
combined "computational neuropsychological" approach can be used. Let us begin
by very briefly considering the key aspects of the Kosslyn & Shwartz
computational theory of visual mental imagery, and then consider one example of
1) how available data in the neuropsychological literature bears on the theory;
2) how behavioral dysfunction following brain damage can be used to test and
help develop the theory; and 3) how PET scanning studies can be used to test
the theory.

The key claims of the Kosslyn & Shwartz theory can be divided into two
classes, pertaining to representations and processes. With regard to
representations, the theory claims that the experience of "having an image"
reflects the existence of a depictive representation in a visual short-term
memory buffer. Such a representation depicts in the same way that a pattern of
points in an array in a computer can depict an object (see Kosslyn, 1983).
This representation occurs in a buffer (which is a component of the functional
architecture) that functions as an array, with patterns within it comprising
the image itself. The image is a temporary representation, which is created on
the basis of information stored in long-term memory. We claim that visual
memories of objects are stored in long-term memory in terms of a) perceptual

memories, organized into "chunks" which correspond to parts of objects (e.g., a dog's body, legs, etc might be stored in distinct units) and b) descriptions, which indicate how the chunks are arranged.

With regard to the processing modules themselves, which make use of the representations, let us consider here only those used in generating an image (i.e., creating a short-term memory representation on the basis of information stored in long-term memory). Previous research has suggested that image generation is not a single computation. Rather, generation seems to involve a processing module that actually activates stored perceptual information (called PICTURE in our theory), another that "looks" for locations where other parts belong on partially completed images (called FIND in our theory), and yet a third (called PUT in our theory) that uses descriptions (e.g., "a cushion is flush on a chair's seat") to position additional parts into an imaged object (see Kosslyn, 1981, for a brief overview). For example, in imaging a chair the PICTURE processing module would activate the main form of the chair (called a "skeletal image" in our theory), and in order to image the cushion on the seat the FIND processing module would locate the seat, and the PUT processing module would use the location information (plus its "understanding" of the meaning of the relation "flush on") to provide input to the PICTURE module so that the cushion would be imaged at the correct position relative to the seat. The PUT processing module is putatively responsible for looking up the description of the part and its relation, and uses this information to invoke the FIND and PICTURE modules appropriately.

This theory is based on computational and empirical arguments: On the computational side, the creative properties of image generation (e.g., as

involved in creating a scene from previously isolated elements, such as Ronald Reagan shaking hands with George Washington)--which are useful in visual reasoning--demand some process that coordinates separately stored encodings. And if images can be formed at different sizes and locations, then new parts must be imaged relative to previously placed ones (not relative to some absolute coordinates), which requires "finding" the parts of previously imaged portions of an object before positioning new portions. On the empirical side, it has been found that the ease of forming an image depends in part on the "discriminabilty" of the location at which it is to be put on an imaged object. This result supports the idea that one inspects a partially completed image in the act of integrating in new parts (see Farah & Kosslyn, 1981; Kosslyn, Reiser, Farah & Fliegel, 1983). In addition, findings that people can use descriptions to arrange items into an imaged scene forces one to posit some computation(s) that use descriptions to position segments of an image (see Kosslyn, 1980; Kosslyn et al, 1983).

It is possible, however, to argue that the data (which consist of reaction-times collected from normal subjects) reflect task demands or the like. And one could argue on computational grounds that the PUT and PICTURE modules are not distinct, that activation of the stored information is simply one aspect of the PUT module's operation. Hence it is desirable to have stronger data supporting the proposed computational decomposition.

Data in the literature: an example

There is already information in the literature on brain damage that seems to have direct bearing on the nature of the representations and processes used in imagery. These data indicate that specific deficits are general across a

class of tasks, and seem to implicate problems in processing an array-like image of the sort posited by our theory. In particular, Bisiach & Luzzatti (1978) report that two patients with unilateral visual neglect (i.e., they ignored visual input on the left side) also showed corresponding neglect in their images of scenes encoded prior to the stroke. When asked to image a plazza from a particular point of view and describe what they "saw," they mentioned only objects that should have been to their right side; when then asked to image it from the opposite side, these patients reported "seeing" objects that now were on their right--which were the very ones ignored immediately before, when they were "viewing" from the opposite perspective! This phenomenon was also found when subjects imaged a familiar room from different perspectives. In later work, Bisiach, Luzzatti & Perani (1979) used a more objective task and found the same results: these sort of patients neglect half of their mental images. It is of especial interest that the patients lacked meta-knowledge about their malady. They were unaware that they neglect the left side...which puts strain on an attempt to explain the data in terms of "task demands" based on "tacit knowledge" (as was discussed in the section on cognitive psychological approaches).

These data, then, are exactly what one would expect if our theory is correct, and images are array-like spatial representations with parts on the left side. Unfortunately, these patients also had slight "field cuts" on the left side. Thus, we cannot infer from these results whether the "mind's eye" (the tests used by the FIND processing module, in our theory) were selectively ignoring half of the representation, or whether half of the functional array in which images occur was disrupted. However, in principle the matter could be

settled if patients could be located with neglect but no field cuts.

Evidence collected to test the theory: Brain damage

A recent review concluded that there is data suggesting that imagery is localized in the left, right, or both hemispheres; there was no unambiguous evidence for its localization in the right hemisphere, as is assumed in the common wisdom (see Ehrlichman & Barrett, 1983). And in fact, Farah (1983) reviewed the neuropsychological literature and found evidence that different imagery abilities may be localized differently; in particular, she argued that image generation requires mechanisms that occur in the left hemisphere. But even here the story is not so clearcut, with some results contradicting the generalization. However, unlike earlier theories of imagery, ours posits that the act of generating an image requires the operation of three processing modules working in concert. And it need not be the case that all computations involved in exercising a given abilility are localized in the same place (or even nearby) in the brain. Our theory might, then, offer a way to sort out what now is a muddy picture in the neuropsychological literature.

Kosslyn, Holtzman, Gazzaniga, & Farah (1984) have performed a large set of experiments designed to examine the claim that the module that coordinates multiple parts into a single image (the PUT processing module) is distinct fom the PICTURE and FIND modules. We began by testing image generation of letters of the alphabet in the two isolated hemispheres of a split-brain patient. In our first series of experiments we asked the subject to make spatial judgments about letters of the alphabet, deciding whether upper case letters were composed only of straight lines or included some curves. Robert Weber and his colleagues have demonstrated convincingly that normal subjects require imagery

in order to make these judgments from memory (see Kosslyn, 1980, for a review
of this work). We reasoned that most adults have seen so many letters that if
asked to image one, they do not image a specific letter they once saw (e.g., on
page 43, line 5 of yesterday's New York Times). Rather, they use a stored
description of the letter to generate a "prototypical example". For example, a
capital "a" might be stored as "two lines meeting at the top joined half way by
a horizontal line." The PUT processing module would use such a description to
assemble an image using stored images of lines, and hence the
letter-classification task should be very difficult if the PUT module were not
operating effectively.

To test this idea, we flashed a lower case letter into the left or right
visual field, and asked our subject to decide whether or not the upper case
version had any curved lines (pressing one button if it did, another if it did
not). He showed a huge left hemisphere advantage. This was interesting in
part because his left hemisphere showed superior ability at language and
inference, both of which involve serial processing of symbols, and we posit
that the PUT module performs serial symbol manipulation. Various control
conditions were conducted to show that the right-hemisphere deficit was not due
to its failing to understand the instructions, to know the association between
upper and lower case letters, to retain an image, to make the judgment, or to
combine together separate stages of a task. The deficit seemed to be in
generating the image from stored information.

In order to implicate a deficit in the operation of the PUT module per se,
we needed to show a dissociation between this task and other imagery tasks that
putatively do not require this module. Thus, in other experiments we used

stimuli that presumably need not be imaged from a description of parts in order
to perform the task. In one, names of animals were presented to one hemisphere
or the othe- If the named animal was larger than a goat, the subject was to
press one button; if a goat was larger, he was to press the other button. Now
both hemispheres performed essentially perfectly, and there was absolutely no
difference in response times. This task has been shown to require imagery
when the to-be-compared objects are close in size (e.g., goat vs. hog; see
chapter 9 of Kosslyn, 1980). In this case, however, only the global shapes
(the "skeletal images") are necessary, not the parts.

One could argue that the right hemisphere simply has problems in
generating images of letters because they are language-related materials.
Thus, it is of interest that in another task the right hemisphere failed
miserably when given the same names of animals used in the size comparison
task. Now, however, the question was, do the animal's ears protrude above the
top of its skull? If so, the subject pressed one button; if not, he pressed
another. In this task, an image of the ears must be correctly positioned
relative to the head, and it is this positioning-operation that apparently is
difficult in the right hemisphere of this patient.

Thus, the results served to implicate a distinct PUT processing module:
both hemispheres were comparable in their abilities to form and evaluate images
of global shapes, which requires the PICTURE and FIND modules, but the right
hemisphere showed a selective deficit in tasks that should require the PUT
module to perform.

The point is, then, that we can directly test our computational theory by
taking advantage of the idea that one or another computation may be localized

in a cerebral hemisphere in this patient. We recently have been repeating the studies done with split brain patients, now using normal subjects and looking for reaction time differences. It is interesting that we find small but consistent reaction time differences in normal right-handed male subjects that mirror the dramatic effects we found with the split brain subjects. However these effects are so small that they would not be noteworthy in the absence of the neuropsychological findings. Because the neuropsychological effects are almost qualitative, these sorts of results have the potential of supplying strong evidence for or against computational theories.

In addition, this sort of approach may well help untangle the convoluted story of how abilities are (or are not) localized in the brain. For example, we now need to administer image generation tasks that do or do not require integration of parts using descriptions, and discover whether patients having different sorts of lesions have selective difficulties with the different tasks.

Evidence from intact brains: PET scanning

Drawing inferences from research on brain-damaged patients is slightly suspect because the functions in a damaged brain could possible have become organized in ways different from an intact one. Thus, it is useful to obtain convergent measures using an entirely different methodology. The Cornell Medical College and Harvard Psychology groups are just now planning PET scanning studies. The logic here is as follows: To the extent that tasks share similar processing, there should be similar patterns of activation in the brain. Further, if a theory claims that the same processing module is used in two tasks, then we may find (but not necessarily) that the same region or

regions are activated in both cases. If we do not find this, we must discover how context shifts which parts of the brain are involved in which functions. The initial studies we are conducting are very simple: For example, we will ask subjects to listen to names of common objects, and to image the sound of the object (e.g., a train), its visual appearance, or both at the same time. If our theory is correct, parts of visual cortex--but not auditory cortex--should be activated when one forms a visual image but not an auditory one, and vice versa when one forms an auditory image. (After they finish imaging the words--and the PET scanning is over--we will test the subjects' recognition memory for sounds and pictures, expecting to find better memory for items imaged in the modality being tested; this will provide a check that subjects actually behaved as asked.) In addition, we hypothesize that the two systems will operate independently, even when one forms a multi-modal image (which will cause activation of the regions activated when visual or auditory images were formed in isolation). In later experiments we plan to ask subjects to participate in various imagery tasks that putatively share greater or lesser numbers of processing modules, and will examine the similarity and overlap of activation during each task (for an example of how this logic can serve to illuminate the nature of individual differences, see Kosslyn, Brunn, Cave & Wallach, in press).

In this case, then, the theory serves to provide a framework for interpreting very complex neuropsychological data. In addition, the techniques of cognitive psychology allow us to design tasks to test the theory using these sorts of data. The three approaches, from AI, cognitive psychology, and neuropsychology, clearly complement each other.

VI. Conclusions

In summary, the time seems ripe for a marriage of AI, cognitive psychology, and neuropsychology. Each field has built up a considerable dowry, but has also revealed limitations. The marriage seems likely to be mutually beneficial. Whether a combined approach will indeed provide a major leap forward is, of course, something only time will tell. But it would not be surprising if the study of cognition were greatly enhanced by considering the brain. Cognition is, after all, something the brain does.

Footnote

References

Anderson, J. R. Arguments concerning representations for mental imagery.
      Psychological Review, 1978, 85, 249-277.

Benton, A. Spatial thinking in neurological patients: historical aspects. In
      M. Potegal (Ed.), Spatial abilities: Developmental and Physiological
      Foundations. New York: Academic Press, 1982.

Bisiach, E., & Luzzatti, C. Unilateral neglect of representational space.
      Cortex, 1978, 14, 129-133.

Bisiach, E., Luzzatti, C., & Perani, D. Unilateral neglect, representational
      schema and consciousness. Brain, 1979, 102, 609-618.

Costa, L. D., & Vaughan, H. G. Performances of patients with lateralized
      cerebral lesions: Verbal and perceptual tests. Journal of Nervous
      and Mental Disease, 1962, 134, 162-168.

Ehrlichman, H., & Barrett, J. Right hemispheric specialization from mental
      imagery: a review of the evidence. Brain and Cognition, 1983, 2,
      55-76.

Farah, M. J. The neurological basis of mental imagery: a componential
      analysis. Cognition, in press.

Farah, M. J., & Kosslyn, S. M. Strategy and structure in image generation.
      Cognitive Science, 1981, 4, 371-383.

Hayes-Roth, F. Distinguishing theories of representation: a critique of
      Anderson's "Arguments concerning mental imagery." Psychological
      Review, 1979, 86, 376-392.

Jackson, J. H. On the nature of the duality of the brain. Brain, 1915, 38,

    80-103 (originally published in 1874).

Jackson, J. H. Case of large cerebral tumour without optic neuritis and with

    left hemiplegia and imperception. Royal London Opthalmological

    Hospital Reports, 1876, 8, 434-444.

Jolicoeur, P., & Kosslyn, S. M. Coordinate systems of visual representations

    in memory. Cognitive Psychology, 1983, 15, 301-345..

Kosslyn, S. M. Image and Mind. Cambridge, MA: Harvard University Press,

    1980.

Kosslyn, S. M. The medium and the message in mental imagery.

    Psychological Review, 1981, 88, 46-66.

Kosslyn, S. M. Ghosts in the Mind's Machine. New York: W. W. Norton,

    1983.

Kosslyn, S. M. The computer and the brain. Harvard University manuscript,

    1984.

Kosslyn, S. M., Brunn, J., Cave, K. R., & Wallach, R. W. Individual

    differences in mental imagery ability: a computational analysis.

    Cognition, in press.

Kosslyn, S. M., Holtzman, J. D., Farah, M. J., & Gazzaniga, M. S. A

    computational analysis of mental image generation: evidence from

    functional dissociations in split-brain patients. Harvard University

    manuscript, 1984.

Kosslyn, S. M., Pinker, S., Smith, G. E., & Shwartz, S. P. On the

    demystification of mental imagery. Behavioral and Brain Sciences,

    1979, 2, 535-581 (including "peer review" commentaries).

Kosslyn, S. M., Reiser, B. J., Farah, M. J., & Fliegel, S. L. Generating
    visual images: units and relations. Journal of Experimental
    Psychology: General, 1983, 112, 278-303.

Newell, A., & Simon, H. A. Human Problem Solving. Englewood Cliffs, NJ:
    Prentice Hall, 1972.

Marr, D. Vision. San Francisco: W. H. Freeman, 1982.

Moscovitch, M. Information processing and the cerebral hemispheres. In
    Handbook of Behavioral Neurobiology (Vol 2: Neuropsychology). New
    York: Plenum, 1979.

Pylyshyn, Z. W. Validating computational models: a critique of Anderson's
    indeterminacy of representation claim. Psychological Review, 1979,
    86, 383-394.

Pylyshyn, Z. W. The imagery debate: analogue media versus tacit knowledge.
    Psychological Review, 1981, 88, 16-45.

Ratcliff, G. Disturbances of spatial orientation associated with cerebral
    lesions. In M. Potegal (Ed.), Spatial abilities: Developmental
    and Physiological Foundations. New York: Academic Press,
    1982.

Shepard, R. N., & Cooper, L. A. Mental Images and their Transformations.
    Cambridge MA: MIT Press, 1982.

Siegler, R. S. (Ed.), Children's Thinking: What Develops?. Hillsdale, NJ:
    Erlbaum Associates, 1978.

Simon, D. P., & Simon, H. A. Individual differences in solving physics
    problems. In R. S. Siegler (Ed.), Children's Thinking: What
    Develops? Hillsdale, NJ: Erlbaum, 1978.

Springer, S. P., & Deutsch, G. Left Brain, Right Brain. San Francisco,

    CA: W. H. Freeman & Co., 1981.

Sternberg, S. High-speed scanning in human memory. Science, 1966, 153,

    652-654.

Ungerleider, L. G., & Mishkin, M. Two cortical visual systems. In D. J. Ingle,

    M. A. Goodale, and R. J. W. Mansfield (Eds.), Analysis of Visual

    Behavior. MIT Press: Cambridge, MA, 1982.

Warrington, E. K., & Taylor, A. M. The contribution of the right parietal lobe

    to object recognition. Cortex, 1973, 9, 152-164.

Winston, P. H. The Psychology of Computer Vision. New York: McGraw-Hill,

    1975.

Navy

1 Robert Ahlers
  Code N711
  Human Factors Laboratory
  NAVTRAEQUIPCEN
  Orlando, FL 32813

1 Dr. Meryl S. Baker
  Navy Personnel R&D Center
  San Diego, CA 92152

1 Dr. Alvah Bittner
  Naval Biodynamics Laboratory
  New Orleans, LA 70189

1 Dr. Nick Bond
  Office of Naval Research
  Liaison Office, Far East
  APO San Francisco, CA 96503

1 Lt. Alexander Bory
  Applied Psychology
  Measurement Division
  NAMRL
  NAS Pensacola, FL 32508

1 Dr. Robert Breaux
  NAVTRAEQUIPCEN
  Code N-095R
  Orlando, FL 32813

1 Dr. Robert Carroll
  NAVOP 115
  Washington , DC 20370

1 Dr. Susan Chipman
  Code 442PT
  Office of Naval Research
  800 N. Quincy St.
  Arlington, VA 22217

1 Dr. Stanley Collyer
  Office of Naval Technology
  800 N. Quincy Street
  Arlington, VA 22217

1 CDR Mike Curran
  Office of Naval Research
  800 N. Quincy St.
  Code 270
  Arlington, VA 22217

Navy

1 Dr. Charles E. Davis
  Personnel and Training Research
  Office of Naval Research (Code 442PT)
  800 North Quincy Street
  Arlington, VA 22217

1 Dr. Marshall J. Farr
  Director, Psychological Sciences
  Office of Naval Research (Code 442)
  800 North Quincy Street
  Arlington, VA 22217

1 DR. PAT FEDERICO
  Code P13
  NPRDC
  San Diego, CA 92152

1 Dr. Jim Hollan
  Code 14
  Navy Personnel R & D Center
  San Diego, CA 92152

1 Dr. Ed Hutchins
  Navy Personnel R&D Center
  San Diego, CA 92152

1 Dr. Norman J. Kerr
  Chief of Naval Technical Training
  Naval Air Station Memphis (75)
  Millington, TN 38054

1 Dr. William L. Maloy (02)
  Chief of Naval Education and Training
  Naval Air Station
  Pensacola, FL 32508

1 Dr. James McBride
  Navy Personnel R&D Center
  San Diego, CA 92152

1 Dr William Montague
  NPRDC Code 13
  San Diego, CA 92152

1 Office of the Chief of Naval Operations
  Research Development & Studies Branch
  OP 115
  Washington, DC 20350

1 LT Frank C. Petho, MSC, USN (Ph.D)
  CNET (N-432)
  NAS
  Pensacola, FL 32508

Navy

1 Dr. Bernard Risland (01C)
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Carl Ross
CNET-PDCD
Building 90
Great Lakes NTC, IL 60088

1 Dr. Robert S. Smith
Office of Chief of Naval Operations
OP-987H
Washington, DC 20350

1 Dr. Alfred F. Smode
Senior Scientist
Code 7B
Naval Training Equipment Center
Orlando, FL 32813

1 Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Frederick Steinheiser
CNO - OP115
Navy Annex
Arlington, VA 20370

1 Dr. Martin A. Tolcott
Leader, Psychological Sciences Division
Office of Naval Research
800 N. Quincy St.
Arlinsgon, VA 22217

1 Dr. James Tweeddale
Technical Director
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152

1 DR. MARTIN F. WISKOFF
NAVY PERSONNEL R& D CENTER
SAN DIEGO, CA 92152

1 Mr John H. Wolfe
Navy Personnel R&D Center
San Diego, CA 92152

Navy

1 Dr. Donald Woodward
Office of Naval Research (Code 441)
800 North Quincy Street
Arlington, VA 22217

1 Dr. Steven Zornetzer
Associate Director for Life Sciences
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217

Marine Corps   Army

1 Jerry Lehnus   1 Dr. Harold F. O'Neil, Jr.
  CAT Project Office   Director, Training Research Lab
  HQ Marine Corps   Army Research Institute
  Washington   , DC 20380   5001 Eisenhower Avenue
                                            Alexandria, VA 22333

1 Major Frank Yohannan, USMC
  Headquarters, Marine Corps   1 Commander, U.S. Army Research Institute
  (Code MPI-20)   for the Behavioral & Social Sciences
  Washington, DC 20380   ATTN: PERI-BR (Dr. Judith Orasanu)
                                            5001 Eisenhower Avenue
                                            Alexandria, VA 22333

                                          1 Dr. Ray Perez
                                            5001 Eisenhower Avenue
                                            PERI-II
                                            Alexandria, VA 22333

                                          1 Joseph Psotka, Ph.D.
                                            ATTN: PERI-1C
                                            Army Research Institute
                                            5001 Eisenhower Ave.
                                            Alexandria, VA 22333

                                          1 Dr. Robert Sasmor
                                            U. S. Army Research Institute for the
                                            Behavioral and Social Sciences
                                            5001 Eisenhower Avenue
                                            Alexandria, VA 22333

                                          1 Dr. Joyce Shields
                                            Army Research Institute for the
                                             Behavioral and Social Sciences
                                            5001 Eisenhower Avenue
                                            Alexandria, VA 22333

                                          1 Dr. Hilda Wing
                                            Army Research Institute
                                            5001 Eisenhower Ave.
                                            Alexandria, VA 22333

Air Force

1 Dr. Earl A. Alluisi
HQ, AFHRL (AFSC)
Brooks AFB, TX 78235

1 Mr. Raymond E. Christal
AFHRL/MOE
Brooks AFB, TX 78235

1 Dr. Alfred R. Fregly
AFOSR/NL
Bolling AFB, DC 20332

1 Dr. Genevieve Haddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

1 Dr. Patrick Kyllonen
AFHRL/MOE

Brooks AFB, TX 78235

1 Dr. Roger Pennell
Air Force Human Resources Laboratory
Lowry AFB, CO 80230

1 Dr. Malcolm Ree
AFHRL/MP
Brooks AFB, TX 78235

1 Dr. Joseph Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Department of Defense

1 Dr. Craig I. Fields
Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209

1 Dr. W. Steve Sellman
Office of the Assistant Secretary
of Defense (MRA & L)
2B269 The Pentagon
Washington, DC 20301

1 Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209

1 Dr. Robert A. Wisher
OUSDRE (ELS)
The Pentagon, Room 3D129
Washington, DC 20301

Civilian Agencies

1 Dr. Helen J. Christup
Office of Personnel R&D
1900 E St., NW
Office of Personnel Management
Washington, DC 20015

1 Dr. Arthur Melmed
724 Brown
U. S. Dept. of Education
Washington, DC 20208

1 Dr. Andrew R. Molnar
Office of Scientific and Engineering
   Personnel and Education
National Science Foundation
Washington, DC 20550

1 Dr. David Pearl
Room 10-C-09
5600 Fishers Lane
Rockville, MD 20857

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Private Sector

1 Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Isaac Bejar
Educational Testing Service
Princeton, NJ 08450

1 Dr. Pat Carpenter
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. John B. Carroll
409 Elliott Rd.
Chapel Hill, NC 27514

1 Dr. Lynn A. Cooper
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Emanuel Donchin
Department of Psychology
University of Illinois
Champaign, IL 61820

1 Dr. Jeffrey Elman
University of California, San Diego
Department of Linguistics
La Jolla, CA 92093

1 Dr. Susan Embertson
PSYCHOLOGY DEPARTMENT
UNIVERSITY OF KANSAS
Lawrence, KS 66045

1 Dr. Anders Ericsson
Department of Psychology
University of Colorado
Boulder, CO 80309

1 Dr. Dexter Fletcher
University of Oregon
Department of Computer Science
Eugene, OR 97403

1 Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Private Sector

1 Dr. Don Gentner
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093

1 Dr. Robert Glaser
Learning Research & Development Center
University of Pittsburgh
3939 O'Hara Street
PITTSBURGH, PA 15260

1 Dr. Daniel Gopher
Faculty of Industrial Engineering
     & Management
TECHNION
Haifa 32000
ISRAEL

1 Dr. Henry M. Halff
Halff Resources
4918 33rd Road, North
Arlington, VA 22207

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Dr. Marcel Just
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Demetrios Karis
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

1 Dr. Steven W. Keele
Dept. of Psychology
University of Oregon
Eugene, OR 97403

1 Dr. Scott Kelso
Haskins Laboratories, Inc
270 Crown Street
New Haven, CT 06510

1 Dr. Stephen Kosslyn
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Private Sector

1 Dr. Marcy Lansman
The L. L. Thurstone Psychometric
  Laboratory
University of North Carolina
Davie Hall 013A
Chapel Hill, NC 27514

1 Dr. Michael Levine
Department of Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

1 Dr. Don Lyon
P. O. Box 44
Higley   , AZ 85236

1 Dr. Jay McClelland
Department of Psychology
MIT
Cambridge, MA 02139

1 Dr. Allen Munro
Behavioral Technology Laboratories
1845 Elena Ave., Fourth Floor
Redondo Beach, CA 90277

1 Dr. Donald A McraaN
Cognitive Science, C-015
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. James W. Pellegrino
University of California,
Santa Barbara
Dept. of Psychology
Santa Barabara , CA 93106

1 Dr. Mike Posner
Department of Psychology
University of Oregon
Eugene, OR 97403

1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St. NW
Washington, DC 20007

1 Dr. David Rumelhart
Center for Human Information Processing
Univ. of California, San Diego
La Jolla, CA 92093

Private Sector

1 PROF. FUMIKO SAMEJIMA
  DEPT. OF PSYCHOLOGY
  UNIVERSITY OF TENNESSEE
  KNOXVILLE, TN 37916

1 Dr. Arthur Samuel
  Yale University
  Department of Psychology
  Box 11A, Yale Station
  New Haven, CT 06520

1 Dr. Walter Schneider
  Psychology Department
  603 E. Daniel
  Champaign, IL 61820

1 Dr. David Shucard
  Brain Sciences Laboratories
  Department of Pediatrics
  National Jewish Hospital & Research Ctr
  3800 E. Colfax Ave.
  Denver, CO 80206

1 Dr. H. Wallace Sinaiko
  Program Director
  Manpower Research and Advisory Services
  Smithsonian Institution
  801 North Pitt Street
  Alexandria, VA 22314

1 Dr. Kathryn T. Spoehr
  Psychology Department
  Brown University
  Providence, RI 02912

1 Dr. Robert Sternberg
  Dept. of Psychology
  Yale University
  Box 11A, Yale Station
  New Haven, CT 06520

1 Dr. Perry W. Thorndyke
  Perceptronics, Inc.
  545 Middlefield Road, Suite 140
  Menlo Park, CA 94025

1 Dr. Douglas Towne
  Univ. of So. California
  Behavioral Technology Labs
  1845 S. Elena Ave.
  Redondo Beach, CA 90277

Private Sector

1 Dr. Keith T. Wescourt
  Perceptronics, Inc.
  545 Middlefield Road, Suite 140
  Menlo Park, CA 94025

1 Dr. Christopher Wickens
  Department of Psychology
  University of Illinois
  Champaign, IL 61820

END

FILMED

10-84

DTIC